arXiv:1412.0639v1 [cs.DS] 1 Dec 2014

Beating the Generator-Enumeration Bound for Solvable-Group Isomorphism*

David J. Rosenbaum
University of Washington
Department of Computer Science & Engineering
Email: djr@cs.washington.edu

December 1, 2014

Abstract

We consider the isomorphism problem for groups specified by their multiplication tables. Until recently, the best published bound for the worst-case was achieved by the $n^{\log_p n+O(1)}$ generator-enumeration algorithm. In previous work with Fabian Wagner, we showed an $n^{(1/2)\log_p n+O(\log n/\log\log n)}$ time algorithm for testing isomorphism of p-groups by building graphs with degree bounded by $p+O(1)$ that represent composition series for the groups and applying Luks' algorithm for testing isomorphism of bounded degree graphs.

In this work, we extend this improvement to the more general class of solvable groups to obtain an $n^{(1/2)\log_p n+O(\log n/\log\log n)}$ time algorithm. In the case of solvable groups, the composition factors can be large which prevents previous methods from outperforming the generator-enumeration algorithm. Using Hall's theory of Sylow bases, we define a new object that generalizes the notion of a composition series with small factors but exists even when the composition factors are large. By constructing graphs that represent these objects and running Luks' algorithm, we obtain our algorithm for solvable-group isomorphism. We also extend our algorithm to compute canonical forms of solvable groups while retaining the same complexity.

*A preliminary version of this work appeared as a portion of [22].

1 Introduction

We study the *group isomorphism problem* in which we must decide if two finite groups given as Cayley tables are isomorphic.

While efficient algorithms are known for testing isomorphism of various special types of groups [15, 24, 25, 12, 13, 19, 1, 6, 5, 2, 9], the class of p-groups, and hence the more general solvable groups, are conjectured [1, 6, 5] to contain the hard case of the group isomorphism problem. For these classes, the $n^{\log_p n+O(1)}$ *generator-enumeration bound* [7, 15, 17], where p is the smallest prime dividing the order of the group, has been the tightest worst-case result for several decades. Deriving this bound is straightforward: such a group must have a generating set of size at most $\log_p n$; since any isomorphism is defined by the image of any generating set, we can test isomorphism by considering all $n^{\log_p n+O(1)}$ possible images of the set. Obtaining an upper bound of $n^{(1-\epsilon)\log_p n+O(1)}$ where $\epsilon > 0$ for the class of p-groups was therefore a longstanding open problem [14].

In previous work with Wagner [23] (following [26, 22]), this was accomplished by showing a square root speedup over generator enumeration for the class of p-groups[1]. However, this left the problem of obtaining an improvement over the generator-enumeration algorithm for solvable groups unresolved.

In the present paper, we show a similar upper bound for the class of solvable-groups. Our construction is based on the techniques of [23] combined with Hall's theory of Sylow bases [10].

For purposes of comparison, we review the algorithm of [23]. We say that two composition series are isomorphic if there is an isomorphism that sends each subgroup in the first series to the corresponding subgroup in the second series. The algorithm consists of two main steps:

1) an $n^{(1/2)\log_p n+O(1)}$ time Turing reduction from group isomorphism to composition-series isomorphism and
2) an algorithm for testing p-group composition series isomorphism in $n^{O(p)}$ time.

Step (1) follows by bounding the number of composition series. For step (2), we construct rooted trees whose levels represent the factors in the composition series; the multiplication table is then encoded by attaching gadgets to the leaves. Since the orders of the composition factors bound the number of children at the corresponding levels of the tree and each leaf is connected to a constant number of gadgets, the resulting graph has degree at most $p + O(1)$. This yields a polynomial-time Karp reduction from composition-series isomorphism to low-degree graph isomorphism. Combining this with an $n^{O(d)}$ time algorithm [16, 4, 3] for testing isomorphism of graphs of degree at most d yields an $n^{O(p)}$ time algorithm for p-group composition-series isomorphism as claimed in step (2).

Combining steps (1) and (2) yields an $n^{(1/2)\log_p n+O(p)}$ algorithm for p-groups (we will refer to this as the graph-isomorphism component of the p-group algorithm). This algorithm is faster than generator-enumeration when p is small and slower when it is large. (We consider a prime small if it is at most $\alpha = \log n/\log\log n$ and large if it is greater than α.) By choosing between these two algorithms according to the value of p, we obtain an $n^{(1/2)\log_p n+O(\log n/\log\log n)}$ time algorithm; this gives a square root speedup over generator enumeration regardless of the value of p.

Our main result leverages Hall's theory of Sylow bases [10] to extend this algorithm to solvable groups.

[1] Subsequent to [23], James Wilson (personal communication) showed that the algorithm for p-group isomorphism from [18] runs in at most $n^{c\log_p n+O(1)}$ time where $c < 1/4$. However, his analysis has not been published and is limited to p-groups.

Theorem 1.1. *Solvable-group isomorphism is decidable in* $n^{(1/2)\log_p n + O(\log n/\log\log n)}$ *deterministic time.*

The algorithm for solvable groups follows the same framework but is more complicated. The main conceptual challenge is because solvable groups can have composition factors of large order as well as other composition factors of small order. This is problematic since both generator enumeration and the graph-isomorphism based p-group algorithm just described will take roughly $n^{\log_p n}$ time for a group that has many small composition factors and one large composition factor.

In order to overcome this obstacle, we need a way to (in effect) apply the graph-isomorphism component of the p-group algorithm to the part of the group that corresponds to the small prime factors while applying the generator-enumeration algorithm to the part of the group that corresponds to large prime factors. Since these two parts of a solvable group do not form a direct product decomposition, we need a way of actually combining these two algorithms since we cannot separate the group into independent parts and run the algorithms separately.

Wagner [27] gave a method for reducing the degree of the graph by restricting the isomorphism to be fixed on the quotient of G by a subgroup G_i in the composition series. If there is a subgroup G_i in the composition series whose prime divisors are all large, then the number of ways of fixing the isomorphism on the quotient G/G_i is relatively small so we can test isomorphism of the composition series. Thus, we could handle large composition factors if we had a way of moving all the large primes to the top of the composition series.

Since it is not clear that there is always a composition series with all the large primes at the top, we use a different structure. The key idea in our algorithm for solvable-group isomorphism is to use Sylow bases to separate the large and small prime divisors [2] (according to the threshold $\alpha = \log n/\log\log n$) into subgroups P_1 and P_2 of G such that $G = P_1P_2$. We call the pair (P_1, P_2) an *α-decomposition* for G and define it formally later. We also let (Q_1, Q_2) be an α-decomposition for H. The correctness of this step is guaranteed by the following lemma which follows easily from Hall's theorems [10].

Lemma 1.2. *For any α, solvable-group isomorphism is deterministic polynomial-time Turing-reducible to testing isomorphism of α-decompositions of the group.*

We then choose a composition series S_2 for P_2 and a composition series S_2' for Q_2. There is no need to choose composition series for P_1 and Q_1 since we plan to apply Wagner's degree reduction trick to these subgroups. We call the pairs (P_1, S_2) and (Q_1, S_2') *α-composition pairs* for G and H. We say that (P_1, S_2) is isomorphic to (Q_1, S_2') if there is an isomorphism from G to H that restricts to isomorphisms from P_1 to Q_1 and S_2 to S_2'. By enumerating all possible composition series as in the case for p-groups, we can reduce the problem to α-composition pair isomorphism.

Lemma 1.3. *Testing isomorphism of the α-decompositions (P_1, P_2) and (Q_1, Q_2) of the groups G and H is $n^{(1/2)\log_p n + O(1)}$ deterministic time Turing reducible to testing isomorphism of α-composition pairs for (P_1, P_2) and (Q_1, Q_2) where p is the smallest prime dividing the order of the group.*

It remains to show how to test if two α-composition pairs are isomorphic. Solving this problem is the main challenge in generalizing the p-group algorithm to solvable groups. As before, we

[2] We thank Laci Babai for suggesting this simplification. An earlier version of this work broke G into many factors which made it more complicated.

accomplish this by constructing a graph. However, now our graph for G must represent both the decomposition $G = P_1P_2$ and the composition series S_2. We start by constructing a tree; the top of the tree corresponds to the subgroup P_1 while the bottom corresponds to S_2. The degree of the top part of the tree is reduced to a constant using Wagner's trick at the cost of a factor of $n^{\alpha+O(1)}$. Extra gadgets are used to require any isomorphism to respect the decomposition $G = P_1P_2$. The multiplication table is represented by attaching gadgets to the leaves in the same way as before. The result is a graph that has degree $\alpha+O(1)$ and represents the isomorphism class of the α-composition pair (P_1, S_2). Combining with the $n^{O(d)}$ time algorithm [16, 4, 3] for testing isomorphism of graphs of degree at most d completes the proof of Theorem 1.1.

As in the case of p-groups, we extend our algorithm for solvable-group isomorphism to compute canonical forms of solvable groups within the same amount of time[3].

In Section 2, we reduce solvable-group isomorphism to α-decomposition isomorphism and from α-decomposition isomorphism to α-composition pair isomorphism. In Section 3, we present the reduction from α-composition pair isomorphism to low-degree graph isomorphism. In Section 4, we derive our algorithms for solvable-group isomorphism.

2 Reducing solvable-group isomorphism to α-composition pair isomorphism

In this section, we define the notions of α-decompositions and α-composition pairs and show Turing reductions from solvable-group isomorphism to α-decomposition isomorphism and from α-decomposition isomorphism to α-composition isomorphism. The first reduction can be done in polynomial time using Hall's theorems [10] while the second follows by counting the number of composition series.

From now on, we assume for convenience that the groups G and H have the same order; if this is not the case, then G and H are not isomorphic. We let α be a parameter that we will later set to $\log n/\log\log n$. We start with the definition of an α-decomposition.

Definition 2.1. *Let G be a group. An α-decomposition of G is a pair of subgroups (P_1, P_2) such that*

(a) $G = P_1P_2$,
(b) every prime dividing $|P_1|$ is greater than α and
(c) every prime dividing $|P_2|$ is at most α

We say that the α-decompositions (P_1, P_2) and (Q_1, Q_2) for the groups G and H are isomorphic if there is an isomorphism $\phi : G \to H$ such that $\phi[P_i] = Q_i$ for each i. In order to reduce solvable-group isomorphism to α-decomposition isomorphism, we now recall two of Hall's theorems. First, we need to define a Sylow basis.

Definition 2.2. *Let G be a group whose order has the prime factorization $n = \prod_{i=1}^{\ell} p_i^{e_i}$. A Sylow basis for G is a set $\{P_i' \mid 1 \le i \le \ell\}$ where each P_i' is a Sylow p_i-subgroup of G and $P_i'P_j' = P_j'P_i'$ for all i and j.*

[3]In a follow-up paper [21], we show how to combine this canonization algorithm with a general collision detection framework to reduce the 1/2 in the exponent of Theorem 1.1 to 1/4.

In a Sylow basis $\{P'_i \mid 1 \leq i \leq \ell\}$, we will always assume that each P'_i is a Sylow p_i-subgroup of G. We say that the Sylow bases $\{P_i \mid 1 \leq i \leq \ell\}$ of G and $\{Q_i \mid 1 \leq i \leq \ell\}$ of H are isomorphic if there exists an isomorphism $\phi : G \to H$ such that $\phi[P_i] = Q_i$ for all i. It is easy to construct an α-decomposition from a Sylow basis by letting P_1 be the product of the Sylow subgroups that correspond to primes that are *greater* than α and letting P_2 be the product of the Sylow subgroups that correspond to primes that are *less* than α.

The following theorem is useful for proving that the reduction from solvable-group isomorphism to α-decomposition isomorphism takes polynomial time.

Theorem 2.3 (Hall [10], cf. [20])**.** *A group G is solvable if and only if it has a Sylow basis.*

Two Sylow bases $\{P'_i \mid 1 \leq i \leq \ell\}$ and $\{Q'_i \mid 1 \leq i \leq \ell\}$ of G are conjugate if there exists $g \in G$ such that for all i, $P'^g_i = Q'_i$.

Theorem 2.4 (Hall [10], cf. [20])**.** *Any two Sylow bases of a solvable group are conjugate.*

Notice that this implies that the group G has at most n Sylow bases. We also require the ability to compute a Sylow basis of a solvable group. This was shown by Kantor and Taylor [11] in the setting of permutation groups so it also holds in our case where the group is specified by its Cayley table.

Theorem 2.5 (Kantor and Taylor [11])**.** *A Sylow basis of a solvable group can be computed deterministically in polynomial time.*

Armed with these results, it is now easy to reduce solvable-group isomorphism to α-decomposition isomorphism. The following lemma from the introduction explains why our results are restricted to the class of solvable groups.

Lemma 1.2. *For any α, solvable-group isomorphism is deterministic polynomial-time Turing-reducible to testing isomorphism of α-decompositions of the group.*

Proof. Let G and H be solvable groups of order $n = \prod_{i=1}^{\ell} p_i^{e_i}$. We compute a Sylow basis $\{P'_i \mid 1 \leq i \leq \ell\}$ for G. Define $P_1 = \prod_{i:p_i>\alpha} P'_i$ and $P_2 = \prod_{i:p_i\leq\alpha} P'_i$; this is an α-decomposition for G. We compute a Sylow basis $\{Q'_i \mid 1 \leq i \leq \ell\}$ for H and consider all of its n conjugates $\left\{Q'^h_i \;\middle|\; 1 \leq i \leq \ell\right\}$ where $h \in H$. For each of these, we define $Q_1 = \prod_{i:p_i<\alpha} Q'_i$ and $Q_2 = \prod_{i:p_i\leq\alpha} Q'_i$) and test if the α-decompositions (P_1, P_2) and (Q_1, Q_2) are isomorphic. We claim that $G \cong H$ if and only if (P_1, P_2) is isomorphic to one of the (Q_1, Q_2) computed above.

Clearly, if G and H are not isomorphic then no α-decomposition of G is isomorphic to an α-decomposition of H. If $\phi : G \to H$ is an isomorphism, then $\{\phi[P'_i] \mid 1 \leq i \leq \ell\}$ is a Sylow basis for H. By Theorem 2.4, it is equal to some conjugate of $\{Q'_i \mid 1 \leq i \leq \ell\}$. Then

$$(Q_1, Q_2) = \Big(\prod_{i:p_i<\alpha} \phi[P'_i], \prod_{i:p_i\leq\alpha} \phi[P'_i]\Big)$$

is an α-decomposition for H that is isomorphic to (P_1, P_2) and our reduction will test if (P_1, P_2) is isomorphic to (Q_1, Q_2). □

Next, we reduce α-decomposition isomorphism to α-composition pair isomorphism. First, we define the notion of an α-composition pair.

Definition 2.6. *An α-composition pair for an α-decomposition (P_1, P_2) of a solvable group G is a pair (P_1, S_2) where S_2 is a composition series for P_2.*

For convenience, we will sometimes say that (P_1, S_2) is an α-composition pair for G. Let (P_1, S_2) and (Q_1, S_2') be a α-decompositions for G and H. Then (P_1, S_2) and (Q_1, S_2') are isomorphic if there is an isomorphism ϕ from (P_1, P_2) to (Q_1, Q_2) which restricts to an isomorphism[4] from S_2 to S_2'.

The reduction from α-decomposition isomorphism to α-composition pair isomorphism, requires an upper bound on the number of composition series for a group and a way to enumerate all composition series. This was shown in [23]; we repeat it here for the convenience of the reader. The argument is based on a suggestion by Laci Babai; previously, we used a more complex argument to enumerate all composition series within a particular class.

Lemma 2.7 ([23])**.** *Let G be a group. Then the number of composition series for G is at most $n^{(1/2)\log_p n + O(1)}$ where p is the smallest prime dividing the order of the group. Moreover, one can enumerate all composition series for G in $n^{(1/2)\log_p n + O(1)}$ time.*

Proof. We show that one can enumerate a class of chains that contains all maximal chains of subgroups in $n^{\log_p n + O(1)}$ time. Since every maximal chain of subgroups contains at most one composition series as a subchain, this suffices to prove the result.

We start by choosing the first nontrivial subgroup in the series. Each of these is generated by a single element so there are at most n choices. If we have a chain $G_0 = 1 < \cdots < G_k$ of subgroups of G, then the next subgroup in the chain can be chosen in at most $|G/G_k|$ ways since different representatives of the same coset generate the same subgroup. Since each $|G_{i+1}| \geq p\,|G_i|$, we see that the number of choices $|G/G_k|$ for G_{k+1} is at most n/p^k. Therefore, the total number of choices required to construct a chain of subgroups in this manner is at most

$$\begin{aligned}\prod_{k=0}^{\lfloor \log_p n \rfloor - 1} (n/p^k) &\leq p^{\sum_{k=0}^{\lceil \log_p n \rceil} k} \\ &= p^{(1/2)\log_p^2 n + O(\log_p n)} \\ &\leq n^{(1/2)\log_p n + O(1)}\end{aligned}$$

Since the set of subgroup chains enumerated by this process includes all maximal chains of subgroups, the result follows. □

We are now ready to derive the algorithm discussed in the introduction.

Lemma 1.3. *Testing isomorphism of the α-decompositions (P_1, P_2) and (Q_1, Q_2) of the groups G and H is $n^{(1/2)\log_p n + O(1)}$ deterministic time Turing reducible to testing isomorphism of α-composition pairs for (P_1, P_2) and (Q_1, Q_2) where p is the smallest prime dividing the order of the group.*

Proof. Let S_2 be an arbitrary composition series for P_2. For each composition series S_2' for Q_2, we test if the α-composition pairs (P_1, S_2) and (Q_1, S_2') are isomorphic. If $\phi : (P_1, P_2) \to (Q_1, Q_2)$

[4]Two composition series $G_0 = 1 \lhd \cdots \lhd G_m = G$ and $H_0 = 1 \lhd \cdots \lhd H_{m'} = H$ are isomorphic if there exists an isomorphism $\phi : G \to H$ such that each $\phi[G_i] = H_i$.

is an isomorphism, then $(Q_1, \phi[S_2])$ is an α-composition pair for H that is isomorphic to (P_1, S_2). Thus, the α-decompositions (P_1, P_2) and (Q_1, Q_2) are isomorphic if and only if the α-composition pair (P_1, S_2) is isomorphic to (Q_1, S_2') for some composition series S_2' for Q_2. The order of Q_2 is at most n; the smallest prime dividing the order of Q_2 is equal to the smallest prime dividing the order of H by Definition 2.1. The complexity then follows from Lemma 2.7. □

We can also prove Turing reductions from solvable-group canonization to α-decomposition canonization and from α-decomposition canonization to α-composition canonization. For the convenience of the reader, we explicitly define canonical forms of groups, α-decompositions and α-decomposition pairs.

Definition 2.8. *A map* $\mathrm{Can}_{\mathbf{Grp}}$ *is a canonical form for groups if for each group* G, $\mathrm{Can}_{\mathbf{Grp}}(G)$ *is an* $n \times n$ *multiplication table with elements in* $[n]$ *that is isomorphic to* G*, such that, if* G *and* H *are groups,* $G \cong H$ *if and only if* $\mathrm{Can}_{\mathbf{Grp}}(G) = \mathrm{Can}_{\mathbf{Grp}}(H)$.

Definition 2.9. *A map* $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}$ *is a canonical form for* α*-decompositions if for each* α*-decomposition* (P_1, P_2) *of a group* G, $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}(P_1, P_2) = (M, \psi[P_1], \psi[M_2])$ *such that the following hold.*

(a) M *is an* $n \times n$ *matrix with entries in* $[n]$.
(b) M *is the multiplication table for a group that is isomorphic to* G *under the isomorphism* $\psi : G \to [n]$.
(c) If (P_1, P_2) *and* (Q_1, Q_2) *are* α*-decompositions then* $(P_1, P_2) \cong (Q_1, Q_2)$ *if and only if* $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}(P_1, P_2) = \mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}(Q_1, Q_2)$.

Definition 2.10. *A map* $\mathrm{Can}_{\alpha\text{-}\mathbf{Pair}}$ *is a canonical form for* α*-composition pairs if for each* α*-composition pair* $(P_1, S_2 = (P_{2,0} = 1 < \cdots < P_{2,m} = P_2))$ *of an* α*-decomposition* (P_1, P_2) *of a group* G, $\mathrm{Can}_{\alpha\text{-}\mathbf{Pair}}(P_1, S_2) = (M, \psi[P_1], \psi[P_{2,0}], \ldots, \psi[P_{2,m}])$ *such that the following hold.*

(a) M *is an* $n \times n$ *matrix with entries in* $[n]$.
(b) M *is the multiplication table for a group that is isomorphic to* G *under* $\psi : G \to [n]$.
(c) If (P_1, S_2) *and* (Q_1, S_2') *are* α*-decompositions then* $(P_1, S_2) \cong (Q_1, S_2')$ *if and only if* $\mathrm{Can}_{\alpha\text{-}\mathbf{Pair}}(P_1, S_2) = \mathrm{Can}_{\alpha\text{-}\mathbf{Pair}}(Q_1, S_2')$.

Our canonical form reductions now follow via similar techniques.

Lemma 2.11. *Computing the canonical form of a solvable group is polynomial-time Turing reducible to computing canonical forms of* α*-decompositions for the group where* p *is the smallest prime dividing the order of the group.*

Proof. Let G be a solvable group of order $n = \prod_{i=1}^{\ell} p_i^{e_i}$. For each Sylow basis $\{P_i' \mid 1 \le i \le \ell\}$ of G, we let $P_1 = \prod_{i:p_i > \alpha} P_i'$ and $P_2 = \prod_{i:p_i \le \alpha} P_i'$ and compute $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}(P_1, P_2)$. We define $\mathrm{Can}_{\mathbf{Grp}}(G)$ to be the multiplication table of the lexicographically least of these canonical forms. Since two groups are isomorphic if and only if the sets of isomorphism classes of their α-decompositions coincide, it follows that $\mathrm{Can}_{\mathbf{Grp}}$ is a canonical form. By Theorem 2.4, there are at most n Sylow bases for G which can be enumerated in polynomial time. Thus, the reduction can be performed in polynomial time. □

Lemma 2.12. *Computing the canonical form of an* α*-decomposition of a group is* $n^{(1/2)\log_p n + O(1)}$ *time Turing reducible to computing canonical forms of* α*-composition pairs for the group where* p *is the smallest prime dividing the order of the group.*

Proof. Let (P_1, P_2) be an α-decomposition of a group G. We use Lemma 2.7 to enumerate all of the at most $n^{(1/2)\log_p n+O(1)}$ composition series S_2 for P_2. We define $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}(P_1, P_2) = (M, \psi[P_1], \psi[P_{2,m}])$ where $(M, \psi[P_1], \psi[P_{2,0}], \ldots, \psi[P_{2,m}])$ is the lexicographically least canonical form of the α-composition pairs (P_1, S_2) that result from this process. It follows from Definition 2.10 that $\mathrm{Can}_{\alpha\text{-}\mathbf{Decomp}}$ is a canonical form. □

3 α-composition-pair isomorphism and canonization

In this section, we show our reduction from α-composition pair isomorphism to low-degree graph isomorphism. Our reduction also extends to reducing α-composition pair canonization to computing canonical forms of low-degree graphs. Our proofs follow an outline similar to the analogous reduction from composition series isomorphism to low-degree graph isomorphism in the case of p-groups [23], but are more complex due to the more general structure of solvable groups.

3.1 Isomorphism testing

At a high level, our algorithm consists of the following steps. First, we augment our α-composition pair (P_1, P_2) by choosing an ordered generating set $\mathbf{g}$ for the subgroup P_1 (which corresponds to the large primes) to obtain the *augmented α-composition pair* $(P_1, S_2, \mathbf{g})$. We say that a mapping $\phi : G \to H$ is an isomorphism between the augmented α-decompositions $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ for G and H if ϕ is an α-composition pair isomorphism for (P_1, S_2) and (Q_1, S_2') and $\phi(\mathbf{g}) = \mathbf{h}$. The reason for choosing an augmented α-composition pair is so that we can reduce the degree of the part of the graph we construct that corresponds to P_1 using the trick due to Wagner [26] mentioned in the introduction.

Since one can fix an ordered generating set $\mathbf{g}$ for P_1 and consider all possible ordered generating sets for Q_1, it is easy to see that α-composition pair isomorphism is $n^{\log_\alpha n+O(1)}$ Turing-reducible to augmented α-composition pair isomorphism. (Recall that we will later set $\alpha = \log n/\log\log n$ so this is $n^{O(\log n/\log\log n)}$ time and is less than the complexity we are aiming for.) We state this in the following lemma.

Lemma 3.1. *Testing isomorphism of the α-composition pairs (P_1, S_2) and (Q_1, S_2') for the solvable groups G and H is $n^{\log_\alpha n+O(1)}$ deterministic time Turing reducible to testing isomorphism of augmented α-composition pairs for (P_1, S_2) and (Q_1, S_2') where p is the smallest prime dividing the order of the group.*

We then construct a tree whose leaves represent the elements of G; by using the ordered generating set $\mathbf{g}$ chosen above, we are able to ensure that the degree of this tree is at most $\alpha + O(1)$. By augmenting this tree with gadgets that represent the multiplication table of the group, we obtain an object that represents the isomorphism class of the augmented α-composition pair $(P_1, P_2, \mathbf{g})$. The final step of the algorithm is to apply the following result for computing canonical forms of low-degree graphs.

Theorem 3.2 (Babai and Luks [4, 3])**.** *Canonization of colored graphs of degree at most d is in $n^{O(d)}$ time.*

The main challenge compared to p-group isomorphism [23] is dealing with the fact that some of the prime divisors of a solvable group can be small while others may be large. This is the

main reason why the correctness proof is significantly more complex than for p-groups. Since a p-group has exactly one prime divisor, it was possible to handle the cases of small and large primes separately using a graph-isomorphism based p-group algorithm [23] (which is fast when the prime is small) and the generator-enumeration algorithm (which is fast when the prime is large). On the other hand, for solvable groups, it is necessary to design a hybrid algorithm that is fast for both cases simultaneously.

As mentioned above, the first step in the graph construction is to define a tree for an augmented α-composition pair $(P_1, P_2, \mathbf{g})$. We do this by constructing trees T_1 and T_2 whose leaves correspond to the elements of P_1 and P_2. In order to define the part of the tree corresponding to P_1, we need a way to canonically order the elements of a group given an ordered generating set. For completeness, we state and prove the required properties from [23].

Definition 3.3 ([23])**.** *Let G be a group with an ordered generating set $\mathbf{g} = (g_1, \ldots, g_k)$. Define a total order $\prec_{\mathbf{g}}$ on G by $x \prec_{\mathbf{g}} y$ if $w_{\mathbf{g}}(x) \prec w_{\mathbf{g}}(y)$ where each $w_{\mathbf{g}}(x) = (x_1, \ldots, x_j)$ is the first word in $\{g_1, \ldots, g_k\}^*$ under the standard ordering such that $x = x_1 \cdots x_j$.*

Lemma 3.4 ([23])**.** *Let G and H be groups with ordered generating sets $\mathbf{g} = (g_1, \ldots, g_k)$ and $\mathbf{h} = (h_1, \ldots, h_k)$, and let $x, y \in G$. Then*

(a) $\prec_{\mathbf{g}}$ is a total ordering on G.
(b) if $\phi : G \to H$ is an isomorphism such that each $\phi(g_i) = h_i$, then $x \prec_{\mathbf{g}} y$ if and only if $\phi(x) \prec_{\mathbf{h}} \phi(y)$.
(c) we can decide if $x \prec_{\mathbf{g}} y$ in $O(n\,|\mathbf{g}|)$ time.

Proof. Let $S = \{g_1, \ldots, g_k\}$. For part (a), it is clear that $\prec_{\mathbf{g}}$ is a total order since $w_{\mathbf{g}} : G \to S^*$ is clearly injective and the standard ordering on S^* is a total order.

For part (b), consider an isomorphism $\phi : G \to H$ such that each $\phi(g_i) = h_i$. Then if $w_{\mathbf{g}}(x) = (x_1, \ldots, x_j)$, $w_{\mathbf{h}}(\phi(x)) = (\phi(x_1), \ldots, \phi(x_j))$. Thus, $x \prec_{\mathbf{g}} y$ if and only if $w_{\mathbf{g}}(x) \prec w_{\mathbf{g}}(y)$ if and only if $w_{\mathbf{h}}(\phi(x)) \prec w_{\mathbf{h}}(\phi(y))$ if and only if $x \prec_{\mathbf{h}} y$.

For part (c), it suffices to show how to compute $w_{\mathbf{g}}(x)$ in polynomial time. Consider the Cayley graph $\mathrm{Cay}(G, S)$ for the group G with generating set S. Then the word $w_{\mathbf{g}}(x)$ corresponds to the edges in the minimum length path from 1 to x in $\mathrm{Cay}(G, S)$ that comes first lexicographically. We can find this path in $O(n\,|\mathbf{g}|)$ time by visiting the nodes in breadth-first order starting with 1. At the j^{th} stage, we know $w_{\mathbf{g}}(y)$ for all $y \in G$ at a distance of at most j from the root. We then compute $w_{\mathbf{g}}(x)$ for each x at a distance of $j+1$ from the root by selecting the minimal word $w_{\mathbf{g}}(x) : g_{x,y}$ over all edges (x, y) associated with an element $g_{x,y}$ of S. □

Now we can define the tree that corresponds to P_1. We do this by choosing a balanced binary tree whose leaves are elements of P_1. The choice of this tree is arbitrary so long as it depends only on $\prec_{\mathbf{g}}$. The reason for constructing the trees for P_1 and P_2 separately is that this allows us to ensure that the tree for P_1 has only constant degree. Otherwise, it would have degree $\Omega(n)$ for groups divisible by large primes which would result in a very slow algorithm. Later on, we will combine the trees for P_1 and P_2 to obtain a tree whose leaves correspond to elements of G.

Definition 3.5. *Let P_1 be a group with ordered generating set $\mathbf{g} = (g_1, \ldots, g_k)$. To construct the rooted tree $T(P_1, \mathbf{g})$, we create a leaf node for each element of P_1 and color each node by the number that corresponds to its position in the ordering $\prec_{\mathbf{g}}$; we then arrange the nodes on a line from smallest to largest according to their colors. We attach a parent node to each pair of adjacent leaves*

starting with the smallest pair; if $|P_1|$ is odd, we attach a single parent node to the last leaf. We then arrange the parent nodes just generated on a line according to the ordering on their children and add new parent nodes for them in the same way. We continue in this manner until we obtain a single root node from which all the leaves are descended; this yields the tree $T(P_1, \mathbf{g})$.

Next, we define the tree for the S_2 using a definition from [23]. We start by letting P_2 be the root of the tree. We then partition P_2 into the cosets obtained by taking P_2 mod the subgroup before P_2 in S_2. These are the children of the node P_2. We continue this partitioning process until we obtain cosets in $P_2/1$; these correspond to the leaves. We state the definition for general groups, but in our case the groups will always be solvable.

Definition 3.6 ([23])**.** *Let P_2 be a group and consider the composition series S_2 given by the subgroups $P_{2,0} = 1 \triangleleft \cdots \triangleleft P_{2,m} = P_2$. Then $T(S_2)$ is defined to be the rooted tree whose nodes are $\bigcup_i (P_2/P_{2,i})$. The root node is P_2. The leaf nodes are $\{x\} \in P_2/1$ which we identify with the elements $x \in P_2$. For each node $xP_{2,i+1} \in P_2/P_{2,i+1}$, there is an edge to each $yP_{2,i}$ such that $yP_{2,i} \subseteq xP_{2,i+1}$.*

In order to obtain a tree whose leaves correspond to elements of G, we need to combine the trees for P_1 and S_2. For this, we require a variant of the rooted product [8] called a leaf product [23]. Given two rooted trees, their leaf product is obtained by identifying the root node of a copy of the second tree with each leaf node.

Definition 3.7 ([23])**.** *Let T_1 and T_2 be trees rooted at r_1 and r_2. Then the leaf product $T_1 \odot T_2$ is the tree rooted at r_1 with vertex set*

$$V(T_1) \cup \{(x, y) \mid x \in L(T_1) \text{ and } y \in V(T_2) \setminus \{r_2\}\}$$

The set of edges is

$$\begin{aligned} E(T_1) \cup \{(x, (x, y)) \mid x \in L(T_1) \text{ and } (r_2, y) \in E(T_2)\} \\ \cup \{((x, y), (x, z)) \mid x \in L(T_1) \text{ and } (y, z) \in E(T_2) \text{ where } y, z \neq r_2\} \end{aligned}$$

For convenience, we identify the tuples $(x, (y, z))$, $((x, y), z)$ with (x, y, z) in the vertex set so that leaf products are associative. It is also useful to define leaf products of tree isomorphisms and bijections between the leaves of two trees.

Definition 3.8 ([23])**.** *For each $1 \leq i \leq k$, let T_i and T_i' be trees rooted at r_i and r_i' and let $\phi_i : L(T_i) \to L(T_i')$ be a bijection that extends to a unique isomorphism which we denote by $\hat{\phi} : T_i \to T_i'$. Then the leaf product $\bigodot_{i=1}^k \phi_i : \bigodot_{i=1}^k T_i \to \bigodot_{i=1}^k T_i'$ sends each $(x_1, \ldots, x_j)$ to $(\hat{\phi}_1(x_1), \ldots, \hat{\phi}_j(x_j))$ where each $x_i \in L(T_i)$ for $i < j$, $x_j \in V(T_j) \setminus \{r_j\}$ and $j \leq k$.*

It is easy to see that $\bigodot_{i=1}^k \phi_i$ is a well-defined isomorphism from $\bigodot_{i=1}^k T_i$ to $\bigodot_{i=1}^k T_i'$. We are now finally in a position to define the tree for a augmented α-composition pair.

Definition 3.9. *Let $(P_1, S_2, \mathbf{g})$ be an augmented α-composition pair for a solvable group G. We define $T(P_1, S_2, \mathbf{g}) = T(P_1, \mathbf{g}) \odot T(S_2)$.*

As in the case of p-groups, we cannot attach the aforementioned multiplication gadgets directly to the tree $T(P_1, S_2, \mathbf{g})$ because each leaf be attached to n gadgets and would thus have degree

$\Omega(n)$; this would cause our algorithm to be extremely slow. We resolve this by utilizing the leaf product of $T(P_1, S_2, \mathbf{g})$ with itself so that each multiplication gadget is only attached to a constant number of leaves.

The following notation is convenient as it allows us to easily associate elements of G with nodes in the tree $T(P_1, S_2, \mathbf{g})$. Let $* : \{(x_1, x_2) \mid x_i \in P_i\} \to G$ by $*(x_1, x_2) = x_1x_2$ and note that this is a bijection. Similarly, we define $\bullet : \{(x_1, x_2) \mid x_i \in Q_i\} \to H$ by $\bullet(x_1, x_2) = x_1x_2$. We can then represent each $x \in G$ by the node $*^{-1}(x)$ in $T(P_1, S_2, \mathbf{g})$ and attach the gadget for each multiplication rule $xy = z$ to the nodes $*^{-1}(x)$, $*^{-1}(y)$ and $*^{-1}(z)$. We formalize this in the following definition.

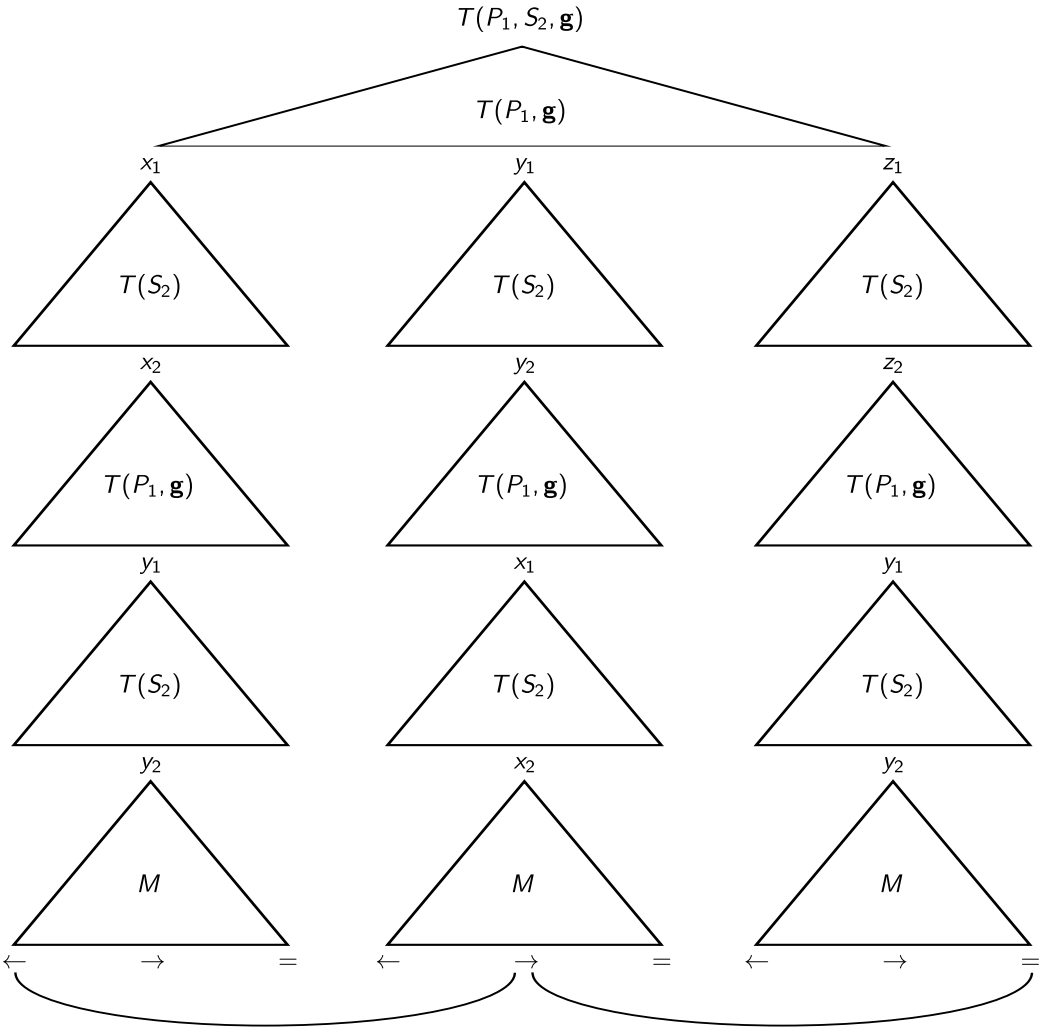

Figure 1: The graph $X(P_1, S_2, \mathbf{g})$ with the multiplication gadget for $xy = z$ where $z = xy$, $*^{-1}(x) = (x_1, x_2)$, $*^{-1}(y) = (y_1, y_2)$ and $*^{-1}(z) = (z_1, z_2)$

Definition 3.10. *Let $(P_1, S_2, \mathbf{g})$ be an augmented α-composition pair for a solvable group G and define M to be the tree with a root connected to three nodes $\leftarrow$, $\rightarrow$ and $=$ with colors "left", "right" and "equals" respectively. We construct $X(P_1, S_2, \mathbf{g})$ by starting with the tree $T(P_1, S_2, \mathbf{g}) \odot$*

$T(P_1, S_2, \mathbf{g}) \odot M$ and connecting multiplication gadgets to the leaf nodes. For each $x, y \in G$, we create the path $((*^{-1}(x), *^{-1}(y), \leftarrow), (*^{-1}(y), *^{-1}(x), \rightarrow), (*^{-1}(xy), *^{-1}(y), =))$. We color each node $(x_1, 1)$ where $x_1 \in P_1$ "second identity." Finally, we color the remaining nodes "internal."

The graph $X(P_1, S_2, \mathbf{g})$ can be thought of a rooted tree with edges added between some nodes at the same levels. The edges from the original tree are called *tree edges* and the edges between nodes at the same level are called *cross edges*. We show $X(P_1, S_2, \mathbf{g})$ in Figure 1.

The correctness of our reduction is based on the fact that two augmented composition pairs $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ are isomorphic if and only if $X(P_1, S_2, \mathbf{g})$ and $X(Q_1, S_2', \mathbf{h})$ are isomorphic. We prove this in the remainder of this subsection.

Some additional terminology is required for the proof. We define **ACP** to be the class of augmented composition pairs for finite solvable groups and let **ACPTree** be the class of graphs that are isomorphic to the graph $X(P_1, S_2, \mathbf{g})$ for some augmented composition pair $(P_1, S_2, \mathbf{g})$. We overload the symbol X from Definition 3.10 by defining $X(\phi) : X(P_1, S_2, \mathbf{g}) \to X(Q_1, S_2', \mathbf{h})$ to be $\phi\big|_{P_1} \odot \phi\big|_{P_2} \odot \phi\big|_{P_1} \odot \phi\big|_{P_2} \odot \mathrm{id}_M$ for each α-composition pair isomorphism $\phi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$.

In order to prove the correctness of our reduction, we need to show that the augmented α-composition pairs $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ are isomorphic if and only if the graphs $X(P_1, S_2, \mathbf{g})$ and $X(Q_1, S_2', \mathbf{h})$ are isomorphic. The forward direction of the implication is equivalent to the assertion that $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})} : \mathrm{Iso}((P_1, S_2, \mathbf{g}), (Q_1, S_2', \mathbf{h})) \to \mathrm{Iso}(X(P_1, S_2, \mathbf{g}), X(Q_1, S_2', \mathbf{h}))$ is well-defined. Proving the converse is more difficult and is one of the main lemmas of this subsection.

Lemma 3.11. *Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H. Then the map*

$$X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})} : \mathrm{Iso}((P_1, S_2, \mathbf{g}), (Q_1, S_2', \mathbf{h})) \to \mathrm{Iso}(X(P_1, S_2, \mathbf{g}), X(Q_1, S_2', \mathbf{h}))$$

is well-defined.

Before proceeding with the proof, it is convenient to introduce additional notation. Let $x, y \in G$. Consider the sequence of nodes that starts at $*^{-1}(x)$, follows tree edges (away from the root) to a node colored "left", follows a cross edge to a node colored "right", then follows tree edges (towards the root) to $*^{-1}(y)$, follows tree edges (away from the root) back to the same node colored "right" and finally follows a cross edge to a node colored "equal"; we call this a *W-sequence* from x to y to xy since its shape resembles a W (see Figure 1). Since W-sequences correspond to multiplication gadgets, there is exactly one W-sequence from $*^{-1}(x)$ to $*^{-1}(y)$: namely, the one that results from the multiplication gadget

$$((*^{-1}(x), *^{-1}(y), \leftarrow), (*^{-1}(y), *^{-1}(x), \rightarrow), (*^{-1}(xy), *^{-1}(y), =)).$$

Therefore, we denote *the* W-sequence from x to y to xy by $W(x, y)$. We now proceed with our proof.

Proof. Consider the augmented α-composition pairs $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ for the solvable groups G and H. Let $\phi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$ be an isomorphism and let $P_{2,0} = 1 \triangleleft \cdots \triangleleft P_{2,m} = P_2$ and $Q_{2,0} = 1 \triangleleft \cdots \triangleleft Q_{2,m} = Q_2$ be the subgroup chains for S_2 and S_2'. Because $\phi(\mathbf{g}) = \mathbf{h}$, it follows from Lemma 3.4 that $\phi\big|_{P_1}$ extends to a unique isomorphism between the rooted colored trees $T(P_1, \mathbf{g})$ and $T(Q_1, \mathbf{h})$. Moreover, since each $\phi[P_{2,i}] = Q_{2,i}$, we see that $\phi\big|_{P_2}$ extends to a unique

isomorphism from $T(S_2)$ to $T(S_2')$. Thus, $\phi\big|_{P_1} \odot \phi\big|_{P_2}$ is an isomorphism from $T(P_1, \mathbf{g}) \odot T(S_2)$ to $T(Q_1, \mathbf{h}) \odot T(S_2')$; therefore, $X(\phi) = \phi\big|_{P_1} \odot \phi\big|_{P_2} \odot \phi\big|_{P_1} \odot \phi\big|_{P_2} \odot \mathrm{id}_M$ is a tree isomorphism.

Let $x, y \in G$ and let $*^{-1}(x) = (x_1, x_2)$. Then $X(\phi)$ maps $*^{-1}(x)$ to $(\phi(x_1), \phi(x_2)) = \bullet^{-1}(\phi(x))$ as $\phi(x) = \phi(x_1)\phi(x_2)$. Similarly, recalling that we identified expressions of the forms $((x_1, x_2), (y_1, y_2))$ and (x_1, x_2, y_1, y_2), we see that $X(\phi)$ maps $(*^{-1}(x), *^{-1}(y))$ to $(\bullet^{-1}(\phi(x)), \bullet^{-1}(\phi(y)))$

Consider the path

$$((*^{-1}(x), *^{-1}(y), \leftarrow), (*^{-1}(y), *^{-1}(x), \rightarrow), (*^{-1}(xy), *^{-1}(y), =))$$

in $X(P_1, S_2, \mathbf{g})$. The image of this path under $X(\phi)$ is

$$((\bullet^{-1}(\phi(x)), \bullet^{-1}(\phi(y)), \leftarrow), (\bullet^{-1}(\phi(y)), \bullet^{-1}(\phi(x)), \rightarrow), (\bullet^{-1}(\phi(xy)), \bullet^{-1}(\phi(y)), =)).$$

By Definition 3.10, this path is one of the multiplication gadgets in $X(Q_1, S_2', \mathbf{h})$. Thus, $X(\phi)$ maps each W-sequence in $X(P_1, S_2, \mathbf{g})$ to a W-sequence in $X(Q_1, S_2', \mathbf{h})$. Moreover, $X(\phi)$ maps each node $(x_1, 1)$ to $(\phi(x_1), 1)$, so it respects the "second identity" color. This implies that $X(P_1, S_2, \mathbf{g}) \cong X(Q_1, S_2', \mathbf{h})$ since both graphs have the same number of multiplication gadgets (and hence the same number of W-sequences). □

In order to show if that if the graphs $X(P_1, S_2, \mathbf{g})$ and $X(Q_1, S_2', \mathbf{h})$ are isomorphic then so are the augmented α-composition pairs $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$, it suffices to show that the map $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})} : \mathrm{Iso}((P_1, S_2, \mathbf{g}), (Q_1, S_2', \mathbf{h})) \to \mathrm{Iso}(X(P_1, S_2, \mathbf{g}), X(Q_1, S_2', \mathbf{h}))$ is surjective. This is the key to our correctness proof and implies that augmented α-composition pair isomorphism reduces to testing isomorphism of the resulting graphs. To do this, we need to show that every isomorphism from $X(P_1, S_2, \mathbf{g})$ to $X(Q_1, S_2', \mathbf{h})$ can be written as a leaf product of group isomorphisms. We accomplish this by restricting the isomorphism between the graphs to certain subsets of nodes and showing that the isomorphism is the leaf product of these restrictions (which turn out to be group isomorphisms). An isomorphism $\theta : X(P_1, S_2, \mathbf{g}) \to X(Q_1, S_2', \mathbf{h})$ induces the bijection $\phi = \bullet \circ \theta \circ *^{-1} : G \to H$. We call this ϕ the *induced bijection* for θ.

Lemma 3.12. *Let $X(P_1, S_2, \mathbf{g})$ and $X(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H, let $\theta : X(P_1, S_2, \mathbf{g}) \to X(Q_1, S_2', \mathbf{h})$ be an isomorphism and let ϕ be its induced bijection. Then*

(a) $\phi : G \to H$ is a group isomorphism,
(b) $\phi_1 = \phi\big|_{P_1} : P_1 \to Q_1$ and $\phi_2 = \phi\big|_{P_2} : P_2 \to Q_2$ are group isomorphisms,
(c) $\theta = \phi_1 \odot \phi_2 \odot \phi_1 \odot \phi_2 \odot \mathrm{id}_M$ and
(d) $\phi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$ is an augmented α-composition pair isomorphism.

Proof. Let us start with part (a). It follows from the assumption that θ is an isomorphism (and hence bijective) that ϕ is a bijection.

Let $x, y \in G$. Now, θ maps the nodes $*^{-1}(x)$ and $*^{-1}(y)$ in $X(P_1, S_2, \mathbf{g})$ to $\bullet^{-1}(\phi(x))$ and $\bullet^{-1}(\phi(y))$ by definition of ϕ. It follows that θ maps the W-sequence $W(x, y)$ from x to y to xy in $X(P_1, S_2, \mathbf{g})$ to the W-sequence $W(\phi(x), \phi(y))$ in $X(Q_1, S_2', \mathbf{h})$. Now, since θ maps $*^{-1}(xy)$ to $\bullet^{-1}(\phi(xy))$, it follows that the W-sequence $W(\phi(x), \phi(y))$ in $X(Q_1, S_2', \mathbf{h})$ is from $\phi(x)$ to $\phi(y)$ to $\phi(xy)$. Therefore, by Definition 3.10, $\phi(xy) = \phi(x)\phi(y)$ so ϕ is a group isomorphism.

Now we prove (b). Let $x_1 \in P_1$. Because θ respects the "second identity" color, it follows that it maps $(x_1, 1)$ to $(x_1', 1)$ for some $x_1' \in Q_1$. Then $x_1' = \phi(x_1)$ which implies that $\phi[P_1] = Q_1$.

Now let $x_2 \in P_2$. Because ϕ is an isomorphism, $\phi(1) = 1$; thus, θ sends the node $(1,1)$ to $(1,1)$ which implies that it maps 1 to 1. Thus, for some $x_2' \in Q_2$,

$$\begin{aligned}\theta(1, x_2) &= (1, x_2') \\ \theta(*^{-1}(x_2)) &= \bullet^{-1}(x_2') \\ \phi(x_2) &= x_2'.\end{aligned}$$

Thus, $\theta(1, x_2) = (1, \phi(x_2))$ so $\phi[P_2] = Q_2$ and ϕ_2 is a group isomorphism.

For part (c), let $x, y \in G$ and $*^{-1}(x) = (x_1, x_2)$. By part (b), θ sends the node x_1 to $\phi_1(x_1)$. Therefore, for some $x_2' \in Q_2$,

$$\begin{aligned}\theta(x_1, x_2) &= (\phi(x_1), x_2') \\ \bullet(\theta(x_1, x_2)) &= \phi(x_1)x_2' \\ \phi(x) &= \phi(x_1)x_2'.\end{aligned}$$

Since $\phi(x) = \phi(x_1)\phi(x_2)$, this implies that $x_2' = \phi(x_2)$ so θ maps $*^{-1}(x) = (x_1, x_2)$ to $\bullet^{-1}(\phi(x)) = (\phi(x_1), \phi(x_2))$.

Now consider a node $(*^{-1}(x), *^{-1}(y), \ell)$ where $x, y \in G$ and $\ell \in \{\leftarrow, \rightarrow, =\}$. As $(*^{-1}(x), *^{-1}(y))$ is in the subtree rooted at $*^{-1}(x)$, θ sends it to a node of the form $(\bullet^{-1}(\phi(x)), \bullet^{-1}(b))$ for some $b \in H$. Similarly, θ maps the node $(*^{-1}(y), *^{-1}(x))$ to a node of the form $(\bullet^{-1}(\phi(y)), \bullet^{-1}(a))$ for some $a \in H$. Now, because $(*^{-1}(x), *^{-1}(y))$ and $(*^{-1}(y), *^{-1}(x))$ are in the W-sequence from x to y to xy, $(\bullet^{-1}(\phi(x)), \bullet^{-1}(b))$ and $(\bullet^{-1}(\phi(y)), \bullet^{-1}(a))$ are in the W-sequence from $\phi(x)$ to $\phi(y)$ to $\phi(xy)$. Then by Definition 3.10, $a = \phi(x)$ and $b = \phi(y)$. Therefore, θ maps $(*^{-1}(x), *^{-1}(y))$ to $(*^{-1}(\phi(x)), *^{-1}(\phi(y)))$. Because of the coloring of the leaves in Definition 3.10, it follows that $\theta = \phi_1 \odot \phi_2 \odot \phi_1 \odot \phi_2 \odot \mathrm{id}_M$.

Finally, let us prove part (d). We already know that ϕ is a group isomorphism by part (a). By part (b), we know that each $\phi[P_i] = Q_i$.

Let $P_{2,0} = 1 \lhd \cdots \lhd P_{2,m} = P_2$ and $Q_{2,0} = 1 \lhd \cdots \lhd Q_{2,m} = Q_2$ be the subgroup chains for S_2 and S_2'. We need to show that each $\phi[P_{2,i}] = Q_{2,i}$. By part (c), θ maps $(1,1)$ in $X(P_1, S_2, \mathbf{g})$ to $(1,1)$ in $X(Q_1, S_2', \mathbf{h})$. Now the path from the root of $X(P_1, S_2, \mathbf{g})$ to $(1,1)$ contains the nodes $(1, P_{2,m}), \ldots, (1, P_{2,0})$ (in that order). Moreover, the descendants of the node $(1, P_{2,i})$ that are in $P_1 \times P_2$ are $\{(1, x_2) \mid x_2 \in P_{2,i}\}$. Similarly, the path from the root of $X(Q_1, S_2', \mathbf{h})$ to $(1,1)$ contains the nodes $(1, Q_{2,m}), \ldots, (1, Q_{2,0})$ (in that order) and the descendants of the node $(1, Q_{2,i})$ that are also in $Q_1 \times Q_2$ are $\{(1, x_2') \mid x_2' \in Q_{2,i}\}$. Therefore, θ maps each set $\{(1, x_2) \mid x_2 \in P_{2,i}\}$ to $\{(1, x_2') \mid x_2' \in Q_{2,i}\}$. Then, by definition of ϕ, $\phi[P_{2,i}] = Q_{2,i}$ and part (d) is proved. □

We now prove that $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is bijective. For isomorphism testing, we only need to show that it is surjective. However, we will need it to be injective later when we discuss canonical forms.

Theorem 3.13. *Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H. Then $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is a bijection. Moreover, both $X(P_1, S_2, \mathbf{g})$ and $X(\phi)$ where $\phi \in \mathrm{Iso}((P_1, S_2, \mathbf{g}), (Q_1, S_2', \mathbf{h}))$ can be computed in polynomial time.*

Proof. We know that $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is well-defined by Lemma 3.11. Let $\theta : X(P_1, S_2, \mathbf{g}) \to X(Q_1, S_2', \mathbf{h})$ be an isomorphism. By Lemma 3.12, the induced bijection $\phi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$

is an isomorphism and $\theta = \phi_1 \odot \phi_2 \odot \phi_1 \odot \phi_2 \odot \mathrm{id}_M$ where each $\phi_i = \phi\big|_{P_i}$. Then $X(\phi) = \theta$ so $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is surjective.

Let $\phi, \psi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$ be isomorphisms and suppose that $X(\phi) = X(\psi)$. Then $\phi_1 \odot \phi_2 \odot \phi_1 \odot \phi_2 \odot \mathrm{id}_M = \psi_1 \odot \psi_2 \odot \psi_1 \odot \psi_2 \odot \mathrm{id}_M$ where each $\phi_i = \phi\big|_{P_i}$ and each $\psi_i = \psi\big|_{P_i}$. Therefore, each $\phi_i = \psi_i$ so $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is injective. □

Correctness of our reduction now follows.

Corollary 3.14. *Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H. Then $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$ if and only if $X(P_1, S_2, \mathbf{g}) \cong X(Q_1, S_2', \mathbf{h})$.*

Because X is defined in terms of leaf products of structures that can be computed in polynomial time, it is immediate that X can also be evaluated in polynomial time.

Lemma 3.15. *Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H and let $\phi : (P_1, S_2, \mathbf{g}) \to (Q_1, S_2', \mathbf{h})$ be an isomorphism. Then both $X(P_1, S_2, \mathbf{g})$ and $X(\phi)$ can be computed in polynomial time.*

The last ingredient that we require for our algorithm for augmented α-composition pair isomorphism is a bound on the degree of the graph.

Lemma 3.16. *Let $(P_1, S_2, \mathbf{g})$ be an augmented α-composition pair for the solvable group G. Then the graph $X(P_1, S_2, \mathbf{g})$ has degree at most $\max\{\alpha + 1, 4\}$ and size $O(n^2)$.*

Proof. The trees $T(P_1, \mathbf{g})$, $T(S_2)$ and M have degrees 3, at most $\alpha + 1$ and 3 respectively. Since $|P_1|\,|P_2| = n$, the size of $T(P_1, \mathbf{g}) \odot T(S_2)$ is $O(n)$. Thus, $T(P_1, \mathbf{g}) \odot T(S_2) \odot T(P_1, \mathbf{g}) \odot T(S_2) \odot M$ has size $O(n^2)$ and degree at most $\max\{\alpha + 1, 4\}$. □

Finally, we obtain our result for augmented α-composition pair isomorphism.

Theorem 3.17. *Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H. Then we can test if $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$ in $n^{O(\alpha)}$ time.*

Proof. By Lemma 3.15, we can compute the graphs $X(P_1, S_2, \mathbf{g})$ and $X(Q_1, S_2', \mathbf{h})$ in polynomial time. By Lemma 3.16 and Theorem 3.2, we can decide if $X(P_1, S_2, \mathbf{g}) \cong X(Q_1, S_2', \mathbf{h})$ in $n^{O(\alpha)}$ time. Finally, Corollary 3.14 tells us that $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$ if and only if $X(P_1, S_2, \mathbf{g}) \cong X(Q_1, S_2', \mathbf{h})$. □

Using Lemma 3.1, we obtain the following corollary.

Corollary 3.18. *Let (P_1, S_2) and (Q_1, S_2') be α-composition pairs for the solvable groups G and H. Then we can test if $(P_1, S_2) \cong (Q_1, S_2')$ in $n^{O(\alpha)+\log_\alpha n}$ time.*

3.2 Canonization

In this subsection, we extend our results for testing isomorphism of α-composition pairs to canonization. This result can be leveraged to obtain faster algorithms for solvable-group isomorphism via collision arguments [21]. Our canonization algorithm requires another map Y that reverses the action of X by sending back to the augmented α-composition pairs from which they arise. We start with the definition for Y. As with X, we overload notation so that Y can also be applied to isomorphisms between graphs.

Definition 3.19. *For each augmented α-composition pair $(P_1, S_2, \mathbf{g})$ for a solvable group G and each graph $A \cong X(P_1, S_2, \mathbf{g})$, we fix an arbitrary isomorphism $\pi : X(P_1, S_2, \mathbf{g}) \to A$. Let $P_{2,0} = 1 \triangleleft \cdots \triangleleft P_{2,m} = P_2$ be the subgroup chain for S_2. Then we define $Y(A) = (\pi[P_1 \times \{1\}], \pi[\{1\} \times P_{2,0}] \triangleleft \cdots \triangleleft \pi[\{1\} \times P_{2,m}], \pi(\mathbf{g}))$.*

Here, $\pi[\{(x_1, x_2) \mid x_i \in P_i\}]$ is interpreted as a group containing each $\pi[\{1\} \times P_{2,i}]$ as a subgroup. For each $x_i, y_i, z_i \in P_i$, we define $\pi(x_1, x_2)\pi(y_1, y_2) = \pi(z_1, z_2)$ if and only if there exists a path $(a_{\pi(x)}a_{\pi(y)}, a_{\pi(z)})$ colored ("left", "right", "equals"), such that $a_{\pi(x)}$, $a_{\pi(y)}$ and $a_{\pi(z)}$ are descendants of the nodes $\pi(x_1, x_2)$, $\pi(y_1, y_2)$ and $\pi(z_1, z_2)$ in the image of the tree $T(P_1, \mathbf{g}) \odot T(S_2) \odot T(P_1, \mathbf{g}) \odot T(S_2) \odot M$ under π.

Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the groups G and H and consider the graphs $A \cong X(P_1, S_2, \mathbf{g})$ and $A' \cong X(Q_1, S_2', \mathbf{h})$. Let $\pi : X(P_1, S_2, \mathbf{g}) \to A$ and $\pi' : X(Q_1, S_2', \mathbf{h}) \to A'$ be the fixed isomorphisms chosen above. Then for each isomorphism $\theta : A \to A'$, we define $Y(\theta) : \pi[\{(x_1, x_2) \mid x_i \in P_i\}] \to \pi'[\{(x_1, x_2) \mid x_i \in Q_i\}]$ to be $\theta\big|_{\pi[\{(x_1,x_2) \mid x_i \in P_i\}]}$.

As for X, we define $Y_{A,A'} : \mathrm{Iso}(A, A') \to \mathrm{Iso}(Y(A), Y(A'))$ by $\theta \mapsto Y(\theta)$ for each pair of graphs $A, A' \in \mathbf{ACPTree}$.

Our first step is to show that Y is well-defined. Once this is proved, we can leverage Theorem 3.13 to show that each $Y_{A,A'}$ is bijective. This allows us to define a canonical form for augmented α-composition pairs in terms of $\mathrm{Can}_{\mathbf{Graph}}$, X and Y.

Lemma 3.20. *Let $(P_1, S_2, \mathbf{g})$ be an augmented α-composition pair for the solvable group G, let A be a graph and let $\pi : X(P_1, S_2, \mathbf{g}) \to A$ be an isomorphism. Then $Y(A)$ is a well-defined augmented composition pair and can be computed in polynomial time. Moreover, $Y(\pi) : (P_1, S_2, \mathbf{g}) \to Y(A)$ is an isomorphism.*

Proof. We claim that $\pi[\{(x_1, x_2) \mid x_i \in P_i\}]$ is indeed a group if interpreted according to Definition 3.19. Let $x_i, y_i, z_i \in P_i$. Then $\pi(x_1, x_2)\pi(y_1, y_2) = \pi(z_1, z_2)$ if and only if there exists a path $(a_{\pi(x)}a_{\pi(y)}, a_{\pi(z)})$ colored ("left", "right", "equals"), such that $a_{\pi(x)}$, $a_{\pi(y)}$ and $a_{\pi(z)}$ are descendants of the nodes $\pi(x_1, x_2)$, $\pi(y_1, y_2)$ and $\pi(z_1, z_2)$ in A. Since π is an isomorphism, this is equivalent to the existence of a path $(a_x a_y, a_z)$ colored ("left", "right", "equals"), such that a_x, a_y and a_z are descendants of the nodes (x_1, x_2), (y_1, y_2) and (z_1, z_2) in $X(P_1, S_2, \mathbf{g})$.

This is in turn equivalent to the existence of a W-sequence from x to y to z where $x = x_1x_2$, $y = y_1y_2$ and $z = z_1z_2$. By definition, this W-sequence exists if and only if $xy = z$. Therefore, $\pi[\{(x_1, x_2) \mid x_i \in P_i\}]$ is a group and $Y(\pi)$ is a group isomorphism from G to $\pi[\{(x_1, x_2) \mid x_i \in P_i\}]$. It is immediate that $Y(A)$ is an augmented α-composition pair and $Y(\pi)$ is an augmented α-composition pair isomorphism.

Now we show how to compute $Y(A)$ in polynomial time. Let $\ell = \lceil \log |P_1| \rceil$ and let the subgroup chain for S_2 be $P_{2,0} = 1 \triangleleft \cdots \triangleleft P_{2,m}$. Then ℓ is the height of $T(P_1, \mathbf{g})$ and m is the height of $T(S_2)$. Thus, by Definition 3.10, $\pi[P_1 \times \{1\}]$ consists of the nodes in A colored "second identity" at a depth of $\ell + m$ from the root.

To compute each $\pi[\{1\} \times P_{2,k}]$, we first find the node $\pi(1, 1)$; this is the identity element of the group $\pi[\{(x_1, x_2) \mid x_i \in P_i\}]$. The node $\pi(1, P_{2,k})$ is the node on the path from the root to $\pi(1, 1)$ in A that is at a distance of $\ell + k$ from the root. Then, by Definition 3.10, each $\pi[\{1\} \times P_{2,k}]$ consists of the nodes in A descended from $\pi(1, P_{2,k})$ that are at a distance of $m - k$ from $\pi(1, P_2)$. □

Now we can show that each $Y_{A,A'}$ is surjective.

Theorem 3.21. *Consider the graphs $A, A' \in \mathbf{ACPTree}$. Then $Y_{A,A'}$ is a bijection and both $Y(A)$ and $Y(\theta)$ where $\theta \in \mathrm{Iso}(Y(A), Y(A'))$ can be computed in polynomial time.*

Proof. Let $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ be augmented α-composition pairs for the solvable groups G and H such that $\pi : X(P_1, S_2, \mathbf{g}) \to A$, $\pi' : X(Q_1, S_2', \mathbf{h}) \to A'$ and $\theta : A \to A'$ are isomorphisms.

First, we observe that Y respects composition and let $\psi = \theta\pi : X(P_1, S_2, \mathbf{g}) \to A'$. Since θ and π are isomorphisms so is ψ; Lemma 3.20 then implies that $Y(\psi) = Y(\theta)Y(\pi)$ is also an isomorphism. Therefore, $Y(\theta) = Y(\psi)(Y(\pi))^{-1}$ is an isomorphism and so $Y_{A,A'}$ is a well-defined function.

Now we prove that $Y_{A,A'}$ is a bijection. It follows from Definitions 3.10 and 3.19 that $YX = I_{\mathbf{ACP}}$. By Theorem 3.13, $X_{(P_1,S_2,\mathbf{g}),(Q_1,S_2',\mathbf{h})}$ is bijective; this implies that $Y_{X(P_1,S_2,\mathbf{g}),X(Q_1,S_2',\mathbf{h})}$ is also bijective since the identity is bijective. Now we just need to show that $Y_{A,A'}$ is bijective. For each isomorphism $\theta : A \to A'$, there exists an isomorphism $\rho : X(P_1, S_2, \mathbf{g}) \to X(Q_1, S_2', \mathbf{h})$ such that $\theta = \pi'\rho\pi^{-1}$. It follows that $Y(\theta) = Y(\pi')Y(\rho)Y(\pi^{-1})$ from which we see that $Y_{A,A'}$ is indeed bijective.

We already showed that $Y(A)$ can be computed in polynomial time in Lemma 3.20 and it follows easily from Definition 3.19 that $Y(\theta)$ can be computed in polynomial time. □

While Theorem 3.21 is enough to obtain our canonization results, we point out that X and Y form a category equivalence when viewed as functors. Moreover, the results of this section can be derived from this more general fact.

To construct our canonical form for augmented α-composition pairs, we convert our augmented α-composition pairs to graphs of degree at most $\alpha + O(1)$ by applying X. Then we compute the canonical form of the resulting graph using Theorem 3.2 and convert it back into an augmented α-composition pair by applying Y. We use $\mathrm{Can}_{\mathbf{Graph}}$ to denote the map from graphs to their canonical forms from Theorem 3.2.

Theorem 3.22. *$Y \circ \mathrm{Can}_{\mathbf{Graph}} \circ X$ is a canonical form for augmented α-composition pairs. Moreover, for any α-composition pair $(P_1, S_2, \mathbf{g})$, we can compute $(Y \circ \mathrm{Can}_{\mathbf{Graph}} \circ X)(P_1, S_2, \mathbf{g})$ in $n^{O(\alpha)}$ time.*

Proof. Consider two α-composition pairs $(P_1, S_2, \mathbf{g})$ and $(Q_1, S_2', \mathbf{h})$ for the solvable groups G and H. By Corollary 3.14, $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$ if and only if

$$X(P_1, S_2, \mathbf{g}) \cong X(Q_1, S_2', \mathbf{h}).$$

Thus, $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$ if and only if

$$\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g})) = \mathrm{Can}_{\mathbf{Graph}}(X(Q_1, S_2', \mathbf{h}))$$

Now, clearly, if $(P_1, S_2, \mathbf{g}) \cong (Q_1, S_2', \mathbf{h})$,

$$Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))) = Y(\mathrm{Can}_{\mathbf{Graph}}(X(Q_1, S_2', \mathbf{h})))$$

On the other hand, if $(P_1, S_2, \mathbf{g}) \ncong (Q_1, S_2', \mathbf{h})$, then

$$\begin{aligned}
\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g})) &\ncong \mathrm{Can}_{\mathbf{Graph}}(X(Q_1, S_2', \mathbf{h})) \\
Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))) &\ncong Y(\mathrm{Can}_{\mathbf{Graph}}(X(Q_1, S_2', \mathbf{h}))) \\
Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))) &\neq Y(\mathrm{Can}_{\mathbf{Graph}}(X(Q_1, S_2', \mathbf{h}))).
\end{aligned}$$

Thus, $Y \circ \mathrm{Can}_{\mathbf{Graph}} \circ X$ is a complete invariant. Also, $X(P_1, S_2, \mathbf{g}) \cong \mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))$ so since $YX = I_{\mathbf{ACP}}$, we have $(P_1, S_2, \mathbf{g}) \cong Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g})))$ by Theorem 3.21. Thus, $Y \circ \mathrm{Can}_{\mathbf{Graph}} \circ X$ is a canonical form.

Lastly, we show that $Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g})))$ can be computed in $n^{O(\alpha)}$ time. By Theorem 3.13, we can compute $X(P_1, S_2, \mathbf{g})$ in polynomial time. By Lemma 3.16 and Theorem 3.2, it takes $n^{O(\alpha)}$ time to compute $\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))$. Finally, by Theorem 3.21, we can compute $Y(\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g})))$ in polynomial time from $\mathrm{Can}_{\mathbf{Graph}}(X(P_1, S_2, \mathbf{g}))$. □

4 Algorithms for solvable-group isomorphism and canonization

Armed with the results of Sections 2 and 3, it is easy to prove Theorem 1.1 as promised in the introduction.

Theorem 1.1. *Solvable-group isomorphism is decidable in $n^{(1/2)\log_p n + O(\log n/\log\log n)}$ deterministic time.*

Proof. Let α be a parameter to be chosen later. By Lemmas 1.2 and 1.3, we can reduce solvable-group isomorphism to α-composition pair isomorphism in $n^{(1/2\log_p n + O(1))}$ where p is the smallest prime dividing the order of the group. Now, by Lemma 3.1, α-composition pair isomorphism reduces to augmented α-composition pair isomorphism in $n^{\log_\alpha n + O(1)}$. Thus, we can reduce solvable-group isomorphism to augmented α-composition pair isomorphism in $n^{(1/2)\log_p n + \log_\alpha n + O(1)}$ time.

Applying Theorem 3.17, we obtain an $n^{(1/2)\log_p n + \log_\alpha n + O(\alpha)}$ time algorithm for solvable-group isomorphism. The optimal choice for α is $\alpha = \log n/\log\log n$. The complexity is then $n^{(1/2)\log_p n + O(\log n/\log\log n)}$ as claimed.

□

Our algorithm for solvable-group canonization follows by a similar argument.

Theorem 4.1. *Solvable-group canonization is in $n^{(1/2)\log_p n + O(\log n/\log\log n)}$ deterministic time.*

Proof. Let α be a parameter to be chosen later. By Lemmas 2.11 and 2.12, we can reduce solvable-group canonization to α-composition pair canonization in $n^{(1/2\log_p n + O(1))}$ where p is the smallest prime dividing the order of the group. α-composition pair canonization in turn reduces to augmented α-composition pair canonization in $n^{\log_\alpha n + O(1)}$ time by enumerating all possible choices of ordered generating sets and choosing the canonical form that comes first lexicographically. Thus, we can reduce solvable-group canonization to augmented α-composition pair canonization in $n^{(1/2)\log_p n + \log_\alpha n + O(1)}$ time.

Applying Theorem 3.22, we obtain an $n^{(1/2)\log_p n + \log_\alpha n + O(\alpha))}$ time algorithm for solvable-group canonization. The optimal choice for α is $\alpha = \log n/\log\log n$. The complexity is again $n^{(1/2)\log_p n + O(\log n/\log\log n)}$ as claimed. □

Acknowledgements

I thank Laci Babai for suggesting the simplification of using the concept of a α-decomposition and other comments, Paul Beame and Aram Harrow for useful discussions and feedback, Joshua Grochow for additional references and the anonymous reviewers for helpful comments. Richard Lipton asked if the techniques used for p-groups could be applied to groups of order $2^a p^b$ where p

is an odd prime. This inspired our more efficient algorithm for solvable-group isomorphism. Part of this work was completed at the Center for Theoretical Physics at the Massachusetts Institute of Technology. This research was funded by the DoD AFOSR through an NDSEG fellowship, by the Simons Foundation through a Simons Award for Graduate Students in Theoretical Computer Science (grant #316172) and by the NSF under grants CCF-0916400 and CCF-1111382.

References

[1] L. Babai, P. Codenotti, J. A. Grochow, and Y. Qiao. Code equivalence and group isomorphism. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1395–1408, 2011.

[2] L. Babai, P. Codenotti, and Y. Qiao. Polynomial-time isomorphism test for groups with no abelian normal subgroups (extended abstract). In *39th International Colloquium on Automata, Languages and Programming*, pages 51–62, 2012.

[3] L. Babai, W. M. Kantor, and E. M. Luks. Computational complexity and the classification of finite simple groups. In *Proceedings of the 24th Annual Symposium on Foundations of Computer Science*, pages 162–171, 1983.

[4] L. Babai and E. M. Luks. Canonical labeling of graphs. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, pages 171–183, 1983.

[5] L. Babai and Y. Qiao. Polynomial-time isomorphism test for groups with abelian Sylow towers. In *29th International Symposium on Theoretical Aspects of Computer Science*, pages 453–464, 2012.

[6] P. Codenotti. *Testing Isomorphism of Combinatorial and Algebraic Structures*. PhD thesis, University of Chicago, 2011.

[7] V. Felsch and J. Neubüser. On a programme for the determination of the automorphism group of a finite group. In *Computational Problems in Abstract Algebra*, pages 59–60, 1970.

[8] C. D. Godsil and B. D. McKay. A new graph product and its spectrum. *Bulletin of the Australian Mathematical Society*, 18(1):21–28, 1978.

[9] J. A. Grochow and Y. Qiao. Algorithms for group isomorphism via group extensions and cohomology. 2013, `1309.1776`.

[10] P. Hall. On the Sylow systems of a soluble group. *Proceedings of the London Mathematical Society*, s2-43(1):316–323, 1938.

[11] W. Kantor and D. Taylor. Polynomial-time versions of Sylow's theorem. *Journal of Algorithms*, 9(1):1–17, 1988.

[12] T. Kavitha. Linear time algorithms for Abelian group isomorphism and related problems. *Journal of Computer and System Sciences*, 73(6):986–996, 2007.

[13] F. Le Gall. Efficient isomorphism testing for a class of group extensions. 2008, `0812.2298`.

[14] R. Lipton. An annoying open problem. Gödel's Lost Letter and $\mathsf{P} = \mathsf{NP}$, 2011.

[15] R. Lipton, L. Snyder, and Y. Zalcstein. *The Complexity of Word and Isomorphism Problems for Finite Groups*. Defense Technical Information Center, 1977.

[16] E. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *Journal of Computer and System Sciences*, 25(1):42–65, 1982.

[17] G. L. Miller. On the $n^{\log n}$ isomorphism technique (a preliminary report). In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, pages 51–58, 1978.

[18] E. A. O'Brien. Isomorphism testing for p-groups. *Journal of Symbolic Computation*, 17(2):133–147, 1994.

[19] Y. Qiao, J. Sarma, and B. Tang. On isomorphism testing of groups with normal Hall subgroups. In *28th International Symposium on Theoretical Aspects of Computer Science*, pages 567–578, 2011.

[20] D. Robinson. *A Course in the Theory of Groups*. Graduate Texts in Mathematics. Springer-Verlag, 1996.

[21] D. J. Rosenbaum. Bidirectional collision detection and faster deterministic isomorphism testing. 2013, `1304.3935`. Submitted to the Symposium on the Theoretical Aspects of Computer Science.

[22] D. J. Rosenbaum. Breaking the $n^{\log n}$ barrier for solvable-group isomorphism. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1054–1073, 2013, `1205.0642`.

[23] D. J. Rosenbaum and F. Wagner. Beating the generator-enumeration bound for p-group isomorphism. 2013, `1312.1755`. Submitted to Theoretical Computer Science.

[24] C. Savage. *An $O(n^2)$ algorithm for Abelian group isomorphism*. Computer Studies Program, North Carolina State University, 1980.

[25] N. Vikas. An $O(n)$ algorithm for Abelian p-group isomorphism and an $O(n \log n)$ algorithm for Abelian group isomorphism. *Journal of Computer and System Sciences*, 53(1):1–9, 1996.

[26] F. Wagner. On the complexity of group isomorphism. *Electronic Colloquium on Computational Complexity*, 2011.

[27] F. Wagner. On the complexity of group isomorphism. *Electronic Colloquium on Computational Complexity*, 2012. Revision 2.